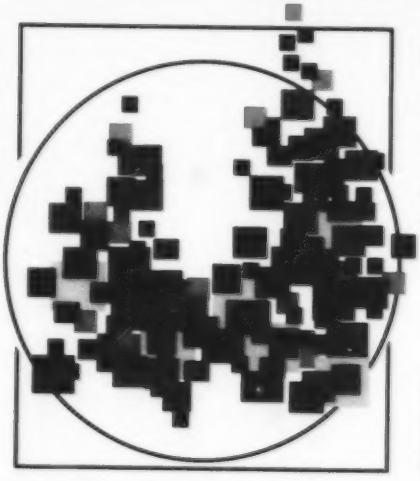

CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Color Constancy from Mutual Reflection

Jian Ho, Mark S. Drew, and Brian V. Funt

CSS/LCCR TR 89-02

Color Constancy from Mutual Reflection

Jian Ho
Mark S. Drew
Brian V. Funt

School of Computing Science
Simon Fraser University
Vancouver, British Columbia, CANADA V5A 1S6
(604) 291-3126

Keywords: Computer vision; color perception; color constancy; finite-dimensional linear model; mutual reflection.

Copyright © 1989 Jian Ho, Mark S. Drew, and Brian V. Funt

92-05566/

Abstract

Mutual reflection occurs when light reflected from one surface illuminates a second surface. In this situation, the color of one or both surfaces can be modified by a "color bleeding" effect. In this paper we examine how sensor values (e.g. RGB values) are modified in the mutual reflection region and show that a good approximation of the surface spectral reflectance function for each surface can be recovered by using the extra information from mutual reflection. Thus, color constancy results from an examination of mutual reflection. Use is made of finite dimensional linear models for ambient illumination and for surface spectral reflectance. If m and n are the number of basis functions required to model illumination and surface spectral reflectance respectively, then we find that the number of different sensor classes p must satisfy the condition $p \geq (2n + m) / 3$. If we use three basis functions to model illumination and three basis functions to model surface spectral reflectance, we find that only three classes of sensors are required to successfully carry out the algorithm. Results are presented showing a very small increase in error over the error inherent in the underlying finite dimension models.

1. Introduction

Two illuminated surfaces of different color can appear to have their colors bleed into one another in regions where light reflected from one surface falls onto the other surface. This is a mutual reflection effect. E.g., in Fig.1 we show two surfaces that form an edge, and the resulting interreflection effect. In this paper we examine how the lights reflected separately from each surface combine to form RGB values in the mutual reflection region. Using vector models for the illumination spectral power distribution and for surface reflectance functions we show that by looking at the extra information coming from measurements of the interreflection RGB values, in addition to sensor measurements from each surface separately, it is possible to recover the surface reflectance properties of both surfaces. In this way we use mutual reflection to obtain "color constancy."

The term *color constancy* refers to the ability of humans and some animals to perceive object colors as approximately constant, independent of changing illumination (see, e.g., [1]). There is evidence that several simultaneous mechanisms contribute to color constancy in humans [2]. Nevertheless, it is the case [3] that an algorithm can be built to recover reflectance from the color signal, in which illumination and and surface characteristics are confounded, without regard to colored surrounds, selective adaptation, memory of colors, etc. Most recent attempts to accomplish this end have made use of finite dimensional linear models in which illumination and surface spectral reflectance are approximated by a weighted sum of a few basis functions of wavelength [4-11].

In the innovative work of Maloney and Wandell [5, 6, 11], several special conditions have to be true for image reflectance to be recoverable from the color signal entering the camera. Some of these conditions are that the illumination must be constant over a given segment of the image, and that a sufficiency of different color information must be available over the image. More importantly, in this method the number of sensor classes (e.g., 3 for RGB sensors) must be at least as great as the number of basis functions modeling the surface spectral reflectance, plus 1. This means that for a usable dimensionality for the set of basis functions modeling reflectance, say 3 or better [12], one must somehow develop a "fourth sensor class" to provide enough information to allow a solution that disentangles illumination and surface spectral reflectance. Alternatively, one can continue using the usual 3 sensor classes, but at the expense of having only a dimensionality of 2 with which to model surface reflectance.

Maloney and Wandell's method essentially involves writing down a set of equations relating the sensor values ρ_k, k=1..p, to the finite dimensional linear model weights, ε_i, i=1..m for illumination and σ_j, j=1..n for surfaces, and then coupling the sets of equations at different pixels by using the assumption of constant illumination. Taking into account enough pixels, that must have sufficiently different sensor values, brings the number of equations up to equal the overall number of unknowns, allowing a solution for ε_i and σ_j.

In [3] it was shown that the need for additional sensor values can essentially be fulfilled by requiring that a finer sampling of the color signal spectrum be introduced as additional input. In that work, it was shown that from detailed knowledge of the spectral power distribution of the color signal a complete disambiguation of the reflectance and illumination from the color signal is possible, and in fact is possible independently at each pixel and without the need for constant illumination or any special requirement on the richness of color variation over image segments. This method [3] relies on a statistical determination of the finite dimensional model weights using a least-squares method, and has the sole requirement that a special mathematical relationship hold amongst the basis

vectors, *viz.* that the set of functions generated by forming products of the basis functions for illumination and surface reflectance form a linearly independent set. The main benefit of this method is that one can increase the dimensionality of the surface reflectance space without having to increase the number of ρ_k values beyond the 3 that correspond to the cones in the human visual system.

In [13, 14] it was shown that the required spectral information for the color signal is derivable from the chromatic aberration effect at edges between adjacent color regions. In fact, what is found via chromatic aberration is the *difference* between color signals from each side of the edge. Coupled with the method of [3] and given this spectral information plus RGB sensor values from each side of the edge, the values of ε_i and σ_j can be determined and color constancy achieved.

In this paper, we take a different approach, although the basic motivation is the same— we seek more information than just the ρ_k values in order to circumvent the "fourth sensor problem." Here, we consider the mutual reflection at a 3-dimensional edge, and the resulting bleeding of colors at the 2-dimensional image edge, without regard to chromatic aberration. Such interreflections generally occur near the edges of objects and the boundaries where objects occlude one another [15].

We can take measurements of ρ_k values from pixels where mutual reflection effects are present and compare them to values where such effects must be absent. By using such measurements on each side of an image edge, relatively far from the edge and also relatively close to the edge, we capture sufficient information to determine ε_i and σ_j separately on each side of the edge, and hence determine the true color for each color patch. This is more than enough information to provide color constancy. In [16] it was shown that mutual reflection can help to distinguish shadow boundaries from material changes. Here we show that mutual reflection has another positive aspect— it is an effect that can be exploited to achieve color constancy. As well, once constant color descriptors are found, the mutual reflection component can be effectively removed from the image. Such an image filter could be useful for shape from shading, etc.

The only constraint on the formulation turns out to be that the number p of sensor classes must be greater than or equal to $(2n + m)/3$, where m and n are the number of basis functions required to model illumination and surface spectral reflectance respectively. This constraint is satisfied by values $p = m = n = 3$ and hence we can keep to only 3 sensor classes while still allowing $n = 3$ for describing reflectances. These values of p, m, and n are shown to recover both surface spectral reflectance and illumination spectral curves well within acceptable errors.

In Section 2 we show how ρ_k measurements on each side of a color edge, and inside the mutual reflection region, are related to the values of ε_i and σ_j for each side of the edge separately, assuming a finite dimensional linear model for describing natural lighting and reflectances. The resulting model is quite complex, but some reasonable assumptions reduce the equations to solvable form. We show that a reciprocity relation exists between the mutual reflection measurements on each side of an edge, so that only ρ_k values in the mutual reflection zone on a single side of the edge are required.

In Section 3 an algorithm is presented for solving for the reflectance weights σ_j on each side of the edge separately, as well as the illumination weights ε_i, which are assumed to not change across an edge.

In Section 4 we examine the accuracy of the solution method. The results are seen to be very good in most cases,

with accuracies of the resulting reflectance spectra not degraded much from the accuracy of the underlying finite dimensional model.

2. Finite-dimensional Models Applied to Mutual Reflection

We assume that illumination and surface spectral reflectance are each modeled to an acceptable degree of accuracy by finite dimensional linear models. Depending on the dimensionality of the set of basis functions, not all illumination spectra and surface reflectances will be well-modeled; generally, the higher the dimensionality of the set of basis functions the better will be the approximation (see, e.g., [12]). However, we shall show below that whatever the accuracy of the underlying finite dimensional model the accuracy of the solution for color constancy from mutual reflection measurements will be not greatly reduced.

For reflectance basis function we use those determined by Cohen via a statistical principal component analysis of 150 Munsell chips randomly selected from a total of 433 chips [17]. For illumination, a similar analysis was carried out by Judd et al. [18] for 622 typical daylight samples. We demonstrate below that it is possible to use $n = m = 3$ basis functions for surfaces and illumination and still have only $p = 3$ sensor classes.

Most naturally occurring daylights and many surface reflectances can be approximated reasonably well using these basis functions, so that by determining a good approximation of the weights ε_i and σ_j, we are in fact deriving an approximation of the entire illumination or reflectance spectrum. Hence what we require from a solution is actually more stringent than simply *color* constancy, in that we seek to recover more spectral information than just chromaticity values. And in fact the measure of accuracy we shall use is the error sum of squares over the entire visible spectrum of an approximate reflectance spectrum compared with an actual one.

Consider the situation depicted in Fig.1. Illumination $E(\lambda)$, assumed constant across an edge, impinges on two surfaces, 1 and 2. Assume surface 1 has spectral reflectance function $S^{(1)}(\lambda)$ and surface 2 has reflectance $S^{(2)}(\lambda)$, all surfaces being ideal diffuse (i.e., Lambertian) reflectors. Then depending on the geometry interreflection may be present as shown in Fig.1. In fact, the two surfaces need not meet at an edge for a mutual reflection effect come into play; and we do not require the surfaces to be flat.

Expanding the illumination function in terms of m basis functions, the spectral power distribution can be approximated by the sum

$$E(\lambda) = \sum_{i=1}^{m} \varepsilon_i E_i(\lambda)$$

with weights ε_i. Similarly, for surface 1 the reflectance can be written

$$S^{(1)}(\lambda) = \sum_{j=1}^{n} \sigma^{(1)}{}_j S_j(\lambda)$$

where reflectance is modeled by n basis functions with n weights $\sigma^{(1)}{}_j$. Surface 2 has reflectance

$$S^{(2)}(\lambda) = \sum_{j=1}^{n} \sigma^{(2)}{}_j S_j(\lambda)$$

with weights $\sigma^{(2)}_j$.

The color signal $I(\lambda)$ is the light reflected by the surface, and is given by the product of E and S. On the surface 1 side, at a point relatively far from the edge the color signal is simply

$$I^{(1)}(\lambda) = E(\lambda)\, S^{(1)}(\lambda).$$

Similarly, on the surface 2 side we have

$$I^{(2)}(\lambda) = E(\lambda)\, S^{(2)}(\lambda).$$

Substituting the basis function decomposition, we have

$$I^{(1)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(1)}_j E_i(\lambda) S_j(\lambda)$$

$$I^{(2)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(2)}_j E_i(\lambda) S_j(\lambda)$$

Now consider the mutual reflection region in which part of $I^{(1)}(\lambda)$, the light reflected from surface 1, impinges onto surface 2. Similarly, part of $I^{(2)}(\lambda)$ impinges onto surface 1. The amount of $I^{(1)}(\lambda)$ that is intercepted by surface 2 is dependent on both the ray geometry between the surfaces and on the shape of each surface, since not every reflected ray is intercepted [consider a small triangle of surface 1 forming an edge with a large plane of surface 2].*

Since there is no *a priori* reason to assume that one surface is more reflective than the other (although for clarity in Fig.1 only a single interreflection is shown), we consider the fraction α_{21} of $I(\lambda)$ from 2 that strikes 1, and the similarly defined α_{12}. Denote by $I_m^{(1)}$ the color signal from a spot on surface 1 in the mutual reflection region, so that it includes a contribution from light reflected by surface 2. It consists of two parts. The first part is $I^{(1)}(\lambda)$, the reflection of $E(\lambda)$ from surface 1. This we assume is the same as from a spot relatively far from the edge, outside the mutual reflection region. The second part is due to the light reflected from surface 2. Since the latter is presumed to come from a spot on surface 2 that is near the edge, the second part is a contribution from $I_m^{(2)}$ (cf. [19]). Hence the color signals coming from each side of the edge in the mutual reflection region are**

$$I_m^{(1)}(\lambda) = I^{(1)}(\lambda) + \alpha_{21}\, I_m^{(2)}(\lambda)\, S^{(1)}(\lambda)$$

$$I_m^{(2)}(\lambda) = I^{(2)}(\lambda) + \alpha_{12}\, I_m^{(1)}(\lambda)\, S^{(2)}(\lambda)$$

where α_{12} is not assumed to be the same as α_{21}.

Solving the above set of simultaneous equations, we can write each $I_m(\lambda)$ in terms of $I^{(1)}(\lambda)$ and $I^{(2)}(\lambda)$ from each surface separately and the geometrical factors α:

---

*The integral of *all* radiance reaching *all* points on the other surface as a fraction of the radiance leaving the first surface into the hemisphere above it is termed the *form-factor* in illumination engineering.

**We include no dependency of the surface reflectance function on angle of incidence or angle of line of sight.

$$I_m^{(1)}(\lambda) = \frac{I^{(1)}(\lambda) + \alpha_{21} I^{(2)}(\lambda) S^{(1)}(\lambda)}{1 - \alpha_{12}\alpha_{21} S^{(1)}(\lambda) S^{(2)}(\lambda)}$$

$$I_m^{(2)}(\lambda) = \frac{I^{(2)}(\lambda) + \alpha_{12} I^{(1)}(\lambda) S^{(2)}(\lambda)}{1 - \alpha_{12}\alpha_{21} S^{(1)}(\lambda) S^{(2)}(\lambda)}$$

The above set of equations would be quite difficult to solve numerically. However, since both α and $S(\lambda)$ are less than 1, not much accuracy is lost by rewriting the equations to first order in $\alpha \cdot S$. This amounts to ignoring more than one bounce of the color signal between the surfaces.

Then we have

$$I_m^{(1)}(\lambda) = I^{(1)}(\lambda) + \alpha_{21} I^{(2)}(\lambda) S^{(1)}(\lambda)$$

$$I_m^{(2)}(\lambda) = I^{(2)}(\lambda) + \alpha_{12} I^{(1)}(\lambda) S^{(2)}(\lambda)$$

as the set to solve.

To see how these equations can be further reduced, denote the terms involving α by

$$\bar{I}_m^{(1)}(\lambda) \equiv I_m^{(1)}(\lambda) - I^{(1)}(\lambda)$$

$$\bar{I}_m^{(2)}(\lambda) \equiv I_m^{(2)}(\lambda) - I^{(2)}(\lambda)$$

That is,

$$\bar{I}_m^{(1)}(\lambda) = \alpha_{21} E(\lambda) S^{(2)}(\lambda) S^{(1)}(\lambda)$$

and also

$$\bar{I}_m^{(2)}(\lambda) = \alpha_{12} E(\lambda) S^{(1)}(\lambda) S^{(2)}(\lambda)$$

Therefore, we have a *reciprocity relation*

$$\bar{I}_m^{(1)}(\lambda) = \frac{\alpha_{21}}{\alpha_{12}} \bar{I}_m^{(2)}(\lambda) .$$

Since all the signals $I(\lambda)$ are in principle measurable, this relation states that one α can be derived from the other. To obtain a solution that takes into account noise and errors in measurement, samples could be measured at N wavelengths so that the above relation would then represent N equations in 1 unknown. A statistical best estimate results from the least-squares solution for the ratio of the α's. Therefore, the number of unknowns to be solved for is reduced by one. Below, we show how it is possible to use only one of the pair of equations for the $I_m(\lambda)$'s to complete a sufficient set to solve for the remaining unknowns.

To simplify further, consider reducing all the the equations involved to equations relating *sensor values* by

multiplying by sensor sensitivity functions $R_k(\lambda)$ for $k = 1 \,..\, p$ sensor classes and integrating over wavelength (e.g., p=3 for RGB signals). Then following [6] and defining the pre-calculated tensor

$$g_{ijk} \equiv \int E_i(\lambda)\, S_j(\lambda)\, R_k(\lambda)\, d\lambda\,,$$

the equations for $I^{(1)}(\lambda)$ and $I^{(2)}(\lambda)$ become the following:

$$\begin{aligned}\rho^{(1)}{}_k &\equiv \int I^{(1)}(\lambda)\, R_k(\lambda)\, d\lambda \\ &= \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(1)}{}_j\, g_{ijk}\,,\end{aligned} \tag{1}$$

and

$$\begin{aligned}\rho^{(2)}{}_k &\equiv \int I^{(2)}(\lambda)\, R_k(\lambda)\, d\lambda \\ &= \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i \sigma^{(2)}{}_j\, g_{ijk}\,,\end{aligned} \tag{2}$$

with $k = 1 \,..\, p$.

For the mutual reflection color signal we first substitute the basis function expansions, yielding

$$I_{\,m}{}^{(2)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \varepsilon_i\, \sigma^{(2)}{}_j\, E_i(\lambda)\, S_j(\lambda) + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \varepsilon_i\, \sigma^{(1)}{}_j\, \sigma^{(2)}{}_{j'}\, E_i(\lambda)\; S_j(\lambda) S_{j'}(\lambda)\,,$$

where we have abbreviated α_{12} simply as α. Now multiplying by $R_k(\lambda)$ and integrating, we find the final equation in our set of equations to solve,

$$\begin{aligned}\rho^{\,m}{}_k &\equiv \int I^{(2)}(\lambda)\, R_k(\lambda)\, d\lambda + \alpha \int I^{(1)}(\lambda)\, S^{(2)}(\lambda)\, R_k(\lambda)\, d\lambda \\ &= \rho^{(2)}{}_k + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \varepsilon_i\, \sigma^{(1)}{}_j\, \sigma^{(2)}{}_{j'}\, h_{ijj'k}\,,\end{aligned} \tag{3}$$

for $k = 1 \,..\, p$, where

$$h_{ijj'k} \equiv \int E_i(\lambda)\, S_j(\lambda)\, S_{j'}(\lambda)\, R_k(\lambda)\, d\lambda$$

and we have abbreviated $\rho^{\,m\,(2)}{}_k$ as $\rho^{\,m}{}_k$.

From the above equation we can see that it is indeed possible to restrict out attention to measurements of $\rho^{(1)}{}_k$, $\rho^{(2)}{}_k$, and $\rho^{\,m}{}_k \equiv \rho^{\,m\,(2)}{}_k$ on *one* side of the edge, since from the similar equation for $\rho^{\,m\,(1)}{}_k$ we can see that if a solution is found for α_{12} then α_{21} can be determined via a least-square solution involving the $\rho^{\,m\,(1)}{}_k$, $\rho^{\,m\,(2)}{}_k$, $\rho^{(1)}{}_k$, and $\rho^{(2)}{}_k$. So that if a solution for ε_i, $\sigma^{(1)}{}_j$, $\sigma^{(2)}{}_j$, and α_{12} can be found via measurements on one side of an edge in the image then a solution for α_{21} is also available, given measurements $\rho^{\,m\,(1)}{}_k$ in the mutual reflection region on the other side.

Using measurements $\rho^{(1)}{}_k$, $\rho^{(2)}{}_k$, and $\rho^{m}{}_k$, we are faced with solving the set of equations (1), (2), and (3) for the ε_i, for the reflectance weights of both surfaces $\sigma^{(1)}{}_j$ and $\sigma^{(2)}{}_j$, and also for the geometrical factor α.

In order that a solution be possible, there must be at least as many equations as unknowns. Counting components in the set $\{\varepsilon_i, \sigma^{(1)}{}_j, \sigma^{(2)}{}_j\}$ we have $m + 2n$ unknowns, and α adds one more. Equations (1), (2), (3) each provide p equations, for a total of $3p$ equations. Hence a solution is possible only if the condition $3p \geq m + 2n + 1$ holds. Therefore we cannot use $p = m = n = 3$ unless we impose a further constraint. We choose to set $\varepsilon_1 \equiv 1$, so that our solutions yield illumination and surface reflectance spectra only up to a multiplicative constant. I.e., for surfaces we determine reflectance but not brightness.

Now our condition reads $p \geq (m + 2n)/3$ and we have a nonlinear set of equations for ε_i, $\sigma^{(1)}{}_j$, $\sigma^{(2)}{}_j$, and α that is at least sufficient for a solution. Since the set of equations is also clearly independent, we would be confident of having fulfilled necessary and sufficient conditions for a solution if the set of equations were linear; however, because they are nonlinear we must proceed with caution.

3. Implementation and Results

We reduce the solution of the nonlinear set of equations (1), (2), (3) to the solution of *linear* ones by breaking up the solution into a multi-stage algorithm. Starting with some reasonable initialization for ε_i, we iteratively solve in turn for $\sigma^{(1)}{}_j$, then $\sigma^{(2)}{}_j$, and finally ε_i. The condition $\varepsilon_1 \equiv 1$ at each step of the iteration also forces α to converge.

In detail, with respect to equations (1), (2), (3), we use the following algorithm:

Initialize $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = 1$

Step 1: Use (1) to solve for $\sigma^{(1)}{}_j$ in terms of ε_i, since (1) is linear in these unknowns.

Step 2: Use (3) to solve for $\alpha \cdot \sigma^{(2)}{}_j$ in terms of ε_i and $\sigma^{(1)}{}_j$.

Step 3: Use (2) to solve for $\varepsilon_i + \alpha$ in terms of $\alpha \cdot \sigma^{(2)}{}_j$.

Step 4: Now set $\varepsilon_1 \equiv 1$: i.e., set $\varepsilon_1 + \alpha \leftarrow 1/\alpha$.

Therefore, set $\alpha \leftarrow \alpha / \varepsilon_1$

and $\varepsilon_i \leftarrow \varepsilon_i / \varepsilon_1$.

Iterate until all values in the set $\{\varepsilon_i, \sigma^{(1)}{}_j, \sigma^{(2)}{}_j, \alpha\}$ change less than a prescribed tolerance.

Since the equations are nonlinear, the order of the above steps is important. In the order described they lead in a stable fashion to the desired solution. In another order, there may be no solution or a spurious one.

An important point to note is the way in which α appears in the algorithm. Since it always appears in combination with another variable until the final normalizing step *the accuracy of solution for all the variables, including* α, *is independent of* α. So long as the mutual reflection effect is indeed present, the solution proceeds to find α. No matter what the actual value of α the percent error in α will be constant for a particular illumination and set of reflectances for the two surfaces. As well, the accuracy of the reflectances recovered by this method will be

independent of the particular value of α dictated by geometry. Of course, this is true mathematically but noise will break this statement down somewhat. We investigate below the effect of noise on the solution method. As well, as $\alpha \rightarrow 0$ we expect the method to fail; this would corrspond to glancing angles of interreflected rays.

In the next section we apply the algorithm to several cases involving naturally occurring lights and reflectances.

4. Applications

We test the method by synthesizing color signals from naturally occurring illumination spectra and surface spectral reflectance functions. By passing these synthetic color signals through bandpass filters representing film or other sensor sensitivities, we determine sensor values ρ_k, which then become the input values for our algorithm, along with the proportion α of the light from one surface reaching the second surface. The algorithm is then used to solve for an estimate of the surface spectral reflectance on each side of the edge, the illumination spectrum, and the mutual reflection factor α.

For surfaces we choose any pair of reflectances from the collection of 370 surface spectral reflectance measurements catalogued by Krinov [20]. From [3] we know how the finite dimensional linear model of reflectance performs for these functions, and we start by choosing a pair of reflectances which are typical in terms of the accuracy of the underlying model. We choose Krinov #53 ("heather, dense growth before flowering") and Krinov #54 ("river valley with meadows") as typical curves. For these reflectances, the underlying model with 3 basis functions taken from Cohen gives errors of 11.04% and 8.11%, respectively. Here, we define an error statistic that measures departures from the correct curve over the entire function of wavelength. It is essentially a rescaled Standard Error divided by the sample root mean square in order to make it a fraction rather than an absolute figure. Error is defined as

$$\{ \sum (y - y')^2 / (\sum y^2) \}^{1/2},$$

where y is an actual value and y′ is the estimate.

For illumination, we can use one of the standard daylights tabulated by Judd et al. [18] or any other tabulated daylight, e.g. those measured by Dixon [21]. Dixon included ultraviolet wavelengths in developing a principal component analysis for Australian daylight, but here we truncate to the visible, 400nm - 650nm. The spectrum for the first characteristic vector for daylight in Bendigo, Australia, is shown in Fig.2 along with Krinov reflectance curve #54 and the resulting color signal composed of the product of the two spectra. Suppose that the shape and geometry of the two surfaces leads to a geometrical factor of $\alpha = 0.05$ for mutual reflection, for rays from surface 1 to surface 2.

The algorithm of the previous section applied to the two color signals generated from this illumination spectrum and Krinov reflectances #53 and #54 can be tested by forming sensor values $\rho^{(1)}{}_k$, $\rho^{(2)}{}_k$, and $\rho^{m}{}_k$. To develop these signals, we used sensitivity functions corresponding to Kodak filters #25 (red), #58 (green) and #47B (blue). The algorithm results are quite good: the value of α found by the solution method is 0.0498, so that the ratio of the estimated α to the actual one is 0.997. As pointed out above, the same accuracy for this ratio is found no matter what the actual value of α is.

To develop an error figure for the spectral reflectance function, it is important to recall that the algorithm gives a

solution only up to an overall multiplicative constant. Therefore, the shape of the reflectance curve is derived, but not its absolute scale. To make a fair comparison, we scale by an appropriate factor and then calculate the error statistic. A best match of the derived curve and the actual one is found by solving for the multiplicative factor by minimizing the squared residuals after multiplying by the unknown factor. The normal equation for this minimization problem yields the factor.

As shown in Fig. 3, the spectral curve for reflectance 2 is quite close to that for Krinov #54; the overall error for the curve found by our algorithm is 9.75%. Since the error of the best fit using the underlying finite dimensional model is 8.11%, the ratio of our model error to the finite dimensional model error is only 1.202. The fit for surface 1, which has reflectance Krinov #53, is similar. We show in Fig. 4 the results of the algorithm in recovering the illumination compared to the actual spectrum. Again, the increase in error is quite small — we obtain an error of 4.13% as compared to the finite dimensional model error of 3.97%.

In Table 1, we show the results of running the algorithm on sensor values derived from the same reflectances as above for all five of Judd's standard daylights, corresponding to five different correlated color temperatures, as well as the Dixon spectrum. As can be seen, the results are similar for all illuminations.

Daylight	α	$\frac{\alpha_{EST}}{\alpha_{ACT}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$\frac{R^{(1)}(\lambda)\text{ Error}}{R^{(1)}(\lambda)\text{ Error}_{FDM}}$	$R^{(2)}(\lambda)$ Error (%)	$\frac{R^{(2)}(\lambda)\text{ Error}}{R^{(2)}(\lambda)\text{ Error}_{FDM}}$
4800°K	0.05006	1.0012	0.40	11.34	1.027	9.85	1.214
5500°K	0.05006	1.0011	0.40	11.33	1.027	9.82	1.211
6500°K	0.05007	1.0015	0.43	11.33	1.026	9.79	1.207
7500°K	0.05006	1.0013	0.41	11.32	1.025	9.77	1.205
10000°K	0.05007	1.0014	0.41	11.32	1.025	9.74	1.201
Dixon	0.04983	0.9966	4.13	11.33	1.026	9.75	1.202

Table 1: Errors using typical reflectances: Krinov #53 and #54, α=0.05, and Judd's North American daylights of five different correlated color temperatures as well as Dixon's Australian daylight. The finite dimensional model best curves have errors of 11.04% for Krinov #53 and 8.11% for Krinov #54. Columns show: (1) α recovered by algorithm; (2) ratio of the recovered α to the actual one; (3) error for illumination spectrum found by algorithm; (4) absolute error of reflectance curve found by algorithm for surface 1; (5) error of reflectance curve for surface 1 compared to error in underlying finite dimensional model; (6,7) similarly for reflectance curve for surface 2. For the Dixon case, the finite dimensional model illumination modeling error is 3.97%, so that the ratio $E(\lambda)$ Error / $E(\lambda)$ Error $_{FDM}$ is 1.039.

A more stringent test is carried out by utilizing atypical reflectance curves that are not well modeled by the basis functions. For this test we use Krinov #162 ("grass, young, green") along with Krinov #54. A best fit using the finite dimensional model itself has an error of 26.65% for Krinov #162. The results for the algorithm are again reasonably good, as seen in Table 2.

Daylight	α	$\frac{\alpha_{EST}}{\alpha_{ACT}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$\frac{R^{(1)}(\lambda)\text{ Error}}{R^{(1)}(\lambda)\text{ Error}_{FDM}}$	$R^{(2)}(\lambda)$ Error (%)	$\frac{R^{(2)}(\lambda)\text{ Error}}{R^{(2)}(\lambda)\text{ Error}_{FDM}}$
4800°K	0.05022	1.0044	3.60	32.31	1.214	11.18	1.378
5500°K	0.05021	1.0043	3.59	32.12	1.205	11.09	1.367
6500°K	0.05023	1.0045	3.49	31.95	1.199	11.00	1.356
7500°K	0.05022	1.0043	3.31	31.84	1.195	10.94	1.348
10000°K	0.05022	1.0044	2.98	31.70	1.189	10.85	1.337
Dixon	0.04999	0.9998	5.74	31.97	1.199	10.95	1.349

Table 2: Poorly fit case: reflectances Krinov #162 and #54, α=0.05 — errors for five Judd daylights and for Dixon daylight (best model fit has error 26.65% for Krinov#162 and 8.11% for Krinov#54); α recovered by algorithm; compared to correct α; illumination spectrum error found by algorithm; reflectance curve 1 error; compared to error in underlying finite dimensional model; similarly for reflectance curve 2. For the Dixon case, the finite dimensional model illumination modeling error is 3.97%, so that the ratio $E(\lambda)$ Error / $E(\lambda)$ Error $_{FDM}$ is 1.446.

To take into account noise, we add a random perturbation to each ρ-vector separately; each component independently receives additive noise equal to a percentage of the vector magnitude. For this test we use the Dixon spectrum coupled to reflectances Krinov #53 and Krinov #54. Running the algorithm repeatedly, we arrive at average results, shown in Table 3. Here, the 0% noise case is shown for comparison, and all other figures are averages. For a few very poorly fit noisy signals at the higher noise level, the algorithm is less stable and a greater tolerance must be used for halting iteration. For a large noise component the results are, of course, not good. Nevertheless, the results are reasonably good for low noise, and in that regime the stability of the algorithm is very good.

Noise	α	$\frac{\alpha_{EST}}{\alpha_{ACT}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$R^{(2)}(\lambda)$ Error (%)
0 %	0.04983	0.9966	4.13	11.33	9.75
5 %	0.04385	0.8771	17.13	16.80	15.05
10 %	0.03912	0.7823	38.15	29.350	27.99

Table 3: Noise: reflectances Krinov #53 and #54, Dixon spectrum, α=0.05 — finite dimensional model best fits have $E(\lambda)$ Error$_{FDM}$=3.97%, $R^{(1)}(\lambda)$ Error$_{FDM}$=11.04%, and $R^{(2)}(\lambda)$ Error$_{FDM}$=8.11%. Columns show: α recovered by algorithm; compared to correct α; illumination spectrum error found by algorithm; reflectance curve 1 error; reflectance curve 2 error.

5. Conclusions

The results show that, in situations in which mutual reflection is present, measurements of the sensor values taken separately from the two different interreflecting surfaces, as well from within the mutual reflection zone, can convey sufficient information to allow complete recovery of both surface spectral reflectance functions as well as of ambient illumination. We also recover the factor α, which summarizes information about the relative positions of the two surfaces. The reason why the method works so well is that by measuring sensor values at two different places on the surface, with one set of measurements coming from a location where both functions are confounded, one is essentially solving the "extra sensors" problem by providing double the information contained in a single set of measurements.

Not only is the result seen to be sufficiently robust to potentially eliminate the mutual reflection effect, but it also in fact recovers a good approximation of complete spectral information. Thus the method yields more than enough information to supply constant color descriptors for the object surfaces, independent of ambient illumination, and therefore color constancy is obtained.

The method recovers geometric information characterized by the factor α. This factor is affected by distance from one reflecting surface to another as well by relative orientation. Hence, shape information might be captured by examining how α changes near an edge, or in general where two surfaces interreflect.

In fact, the algorithm still works even if both surfaces have the *same* color, with the same or differing brightness. This case might correspond to a folded piece of uniform material. The recovery of α here gives shape information near the fold. Why the method works in this one-color case is that the reflected light is a different color from the illumination, and the extra measurements of sensor values supply enough information to disentangle the two. The recovery of shape information from the mutual reflection phenomenon will be explored elsewhere.

Clearly, the interreflection model put forward here, that uses basis function decompositions for an accurate whole-spectrum model of reflection, can be extended to the modeling of multi-surface enclosures for graphics applications (cf. [19]). Multiple interreflections could be incorporated at the expense of having a more difficult mathematical system to solve; for enclosures, such multiple bounces may add a significant contribution.

If use were made in our model of broad-band sensor functions such as those in the human visual system, instead of bandpass filters common in video and film, then the results of the algorithm would be somewhat more accurate since a broader sampling of more frequencies results in a better representation of the entire color signal spectrum in terms of the finite dimensional model weights.

References

1. J. Beck, *Surface Color Perception*, Cornell Univ. Press, Ithaca, N.Y., 1972.
2. K.T. Blackwell and G. Buchsbaum, "Quantitative studies of color constancy", *J. Opt. Soc. Am. A.*, Vol. 5 1988, pp. 1772-1780.
3. J. Ho, B.V. Funt, and M.S. Drew, "Disambiguation of illumination and surface reflectance from spectral power distribution of colour signal: theory and applications", Tech. report CSS/LCCR TR 88-18, Simon Fraser University School of Computing Science,1988.
4. G. Buchsbaum, "A spatial processor model for object colour perception", *J. Franklin Inst.*, Vol. 310 1980, pp. 1-26.
5. L.T. Maloney, *Computational Approaches To Color Constancy*, PhD dissertation, Stanford University, 1985, Applied Psychology Laboratory
6. L.T. Maloney and B.A. Wandell, "Color constancy: a method for recovering surface spectral reflectance", *J. Opt. Soc. Am. A*, Vol. 3 1986, pp. 29-33.
7. M.H. Brill and G. West, "Chromatic adaptation and color constancy: a possible dichotomy", *Col. Res. Appl.*, Vol. 11 1986, pp. 196-204.
8. M. D'Zmura, and P. Lennie, "Mechanisms of color constancy", *J. Opt. Soc. Am.*, Vol. 3 No. 10 1965, pp. 1662-1672.
9. A. Yuille, "A method for computing spectral reflectance", *Biol. Cybern.*, Vol. 56 1987, pp. 195-201.
10. R. Gershon, A.D. Jepson, and J.K. Tsotsos, "From [R,G,B] to surface reflectance: computing color constant descriptors in images", *Proceedings of IJCAI conference*, 1987, pp. 755-758.
11. B.A. Wandell, "The synthesis and analysis of color images", *IEEE Trans. Patt. Anal. and Mach. Intell.*, Vol. PAMI-9 1987, pp. 2-13.
12. L.T. Maloney, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters", *J. Opt. Soc. Am. A*, Vol. 3 1986, pp. 1673-1683.
13. B.V. Funt and J. Ho, "Color from black and white", *Proceedings of the Second International Conference on Computer Vision, Tarpon Springs Dec 5-8*, IEEE, 1988, pp. 2-8, and International Journal of Computer Vision, (in press)
14. J. Ho and B.V. Funt, " Color constancy from chromatic aberration", Tech. report CSS/LCCR TR 88-18, Simon Fraser University School of Computing Science,1988.
15. B.K.P. Horn, *Robot Vision*, MIT Press, 1986.
16. R. Gershon, A.D. Jepson, and J.K. Tsotsos, "Ambient illumination and the determination of material changes", *J. Opt. Soc. Am. A*, Vol. 3 1986, pp. 1700-1707.
17. J. Cohen, "Dependency of The Spectral Reflectance Curves of The Munsell Color Chips", *Psychon. Sci.*, Vol. 1 1964, pp. 369-370.
18. D.B. Judd, D.L. MacAdam, and G. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature", *J. Opt. Soc. Am.*, Vol. 54 August 1964, pp. 1031-1040.
19. C.M. Goral, K.E. Torrance, D.P. Greenberg, and B. Battaile, "Modeling the interaction of light between diffuse surfaces", *Computer Graphics*, Vol. 18 1984, pp. 213-222.
20. E.L. Krinov, "Spectral Reflectance Properties of Natural Formations", *Technical Translation TT-439, National Research Council of Canada*, 1947.
21. E.R. Dixon, "Spectral distribution of Australian daylight", *J. Opt. Soc. Am*, Vol. 68 1978, pp. 437-450.

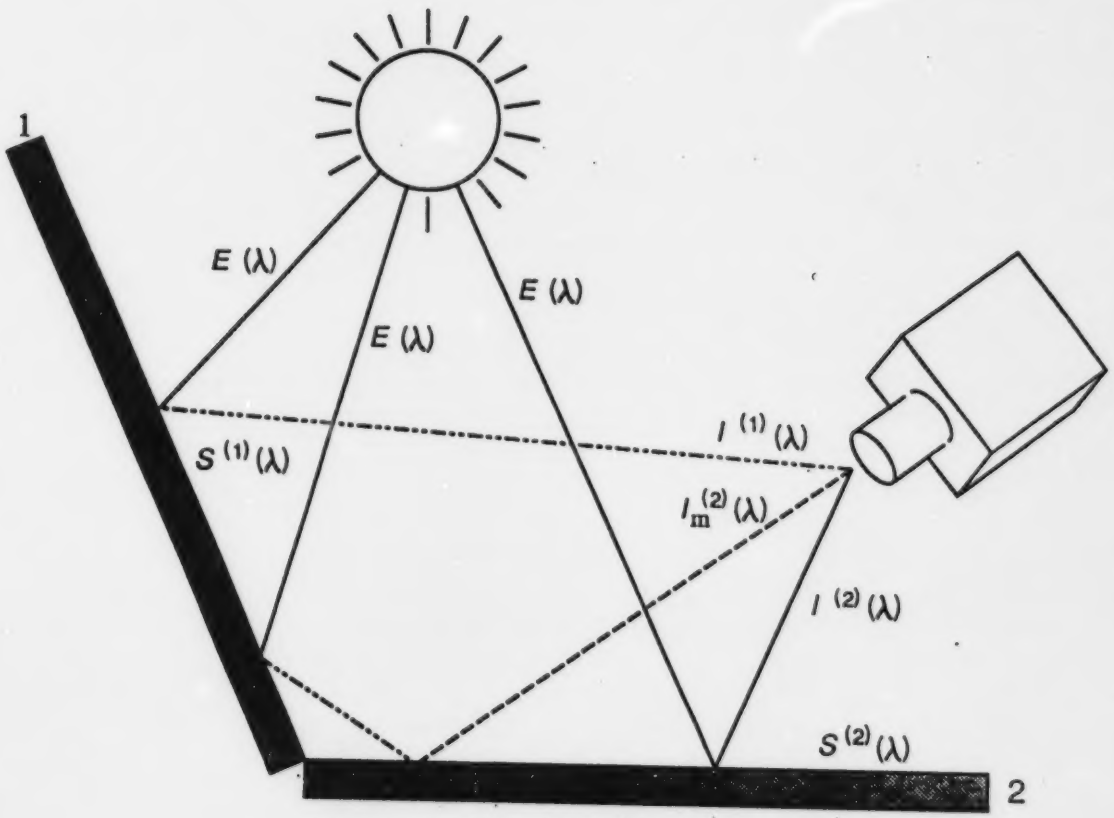

Fig.1. Approximate surface spectral reflectance for surfaces 1 and 2 can be found from color signals from each separately, plus the color signal from the mutual reflection region.

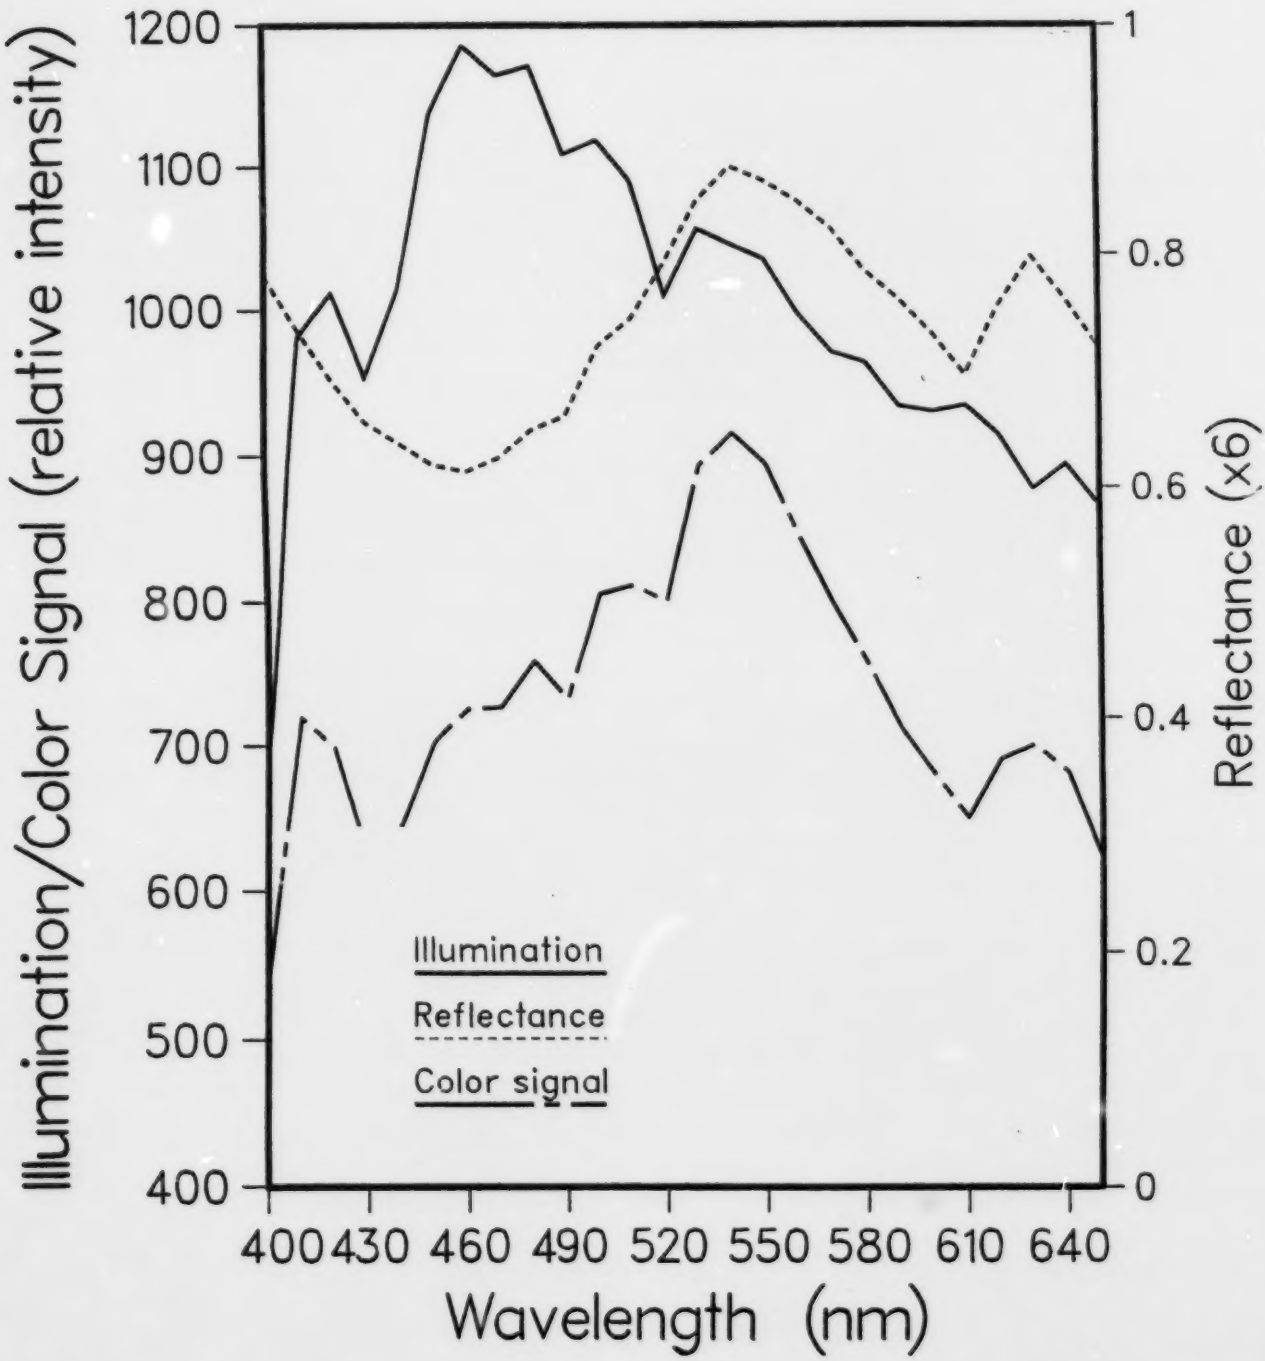

Fig.2. Typical case: color signal formed from product of surface spectral reflectance Krinov #54 and Dixon spectrum for Australian daylight. The color signal has been scaled up by a factor of 6.

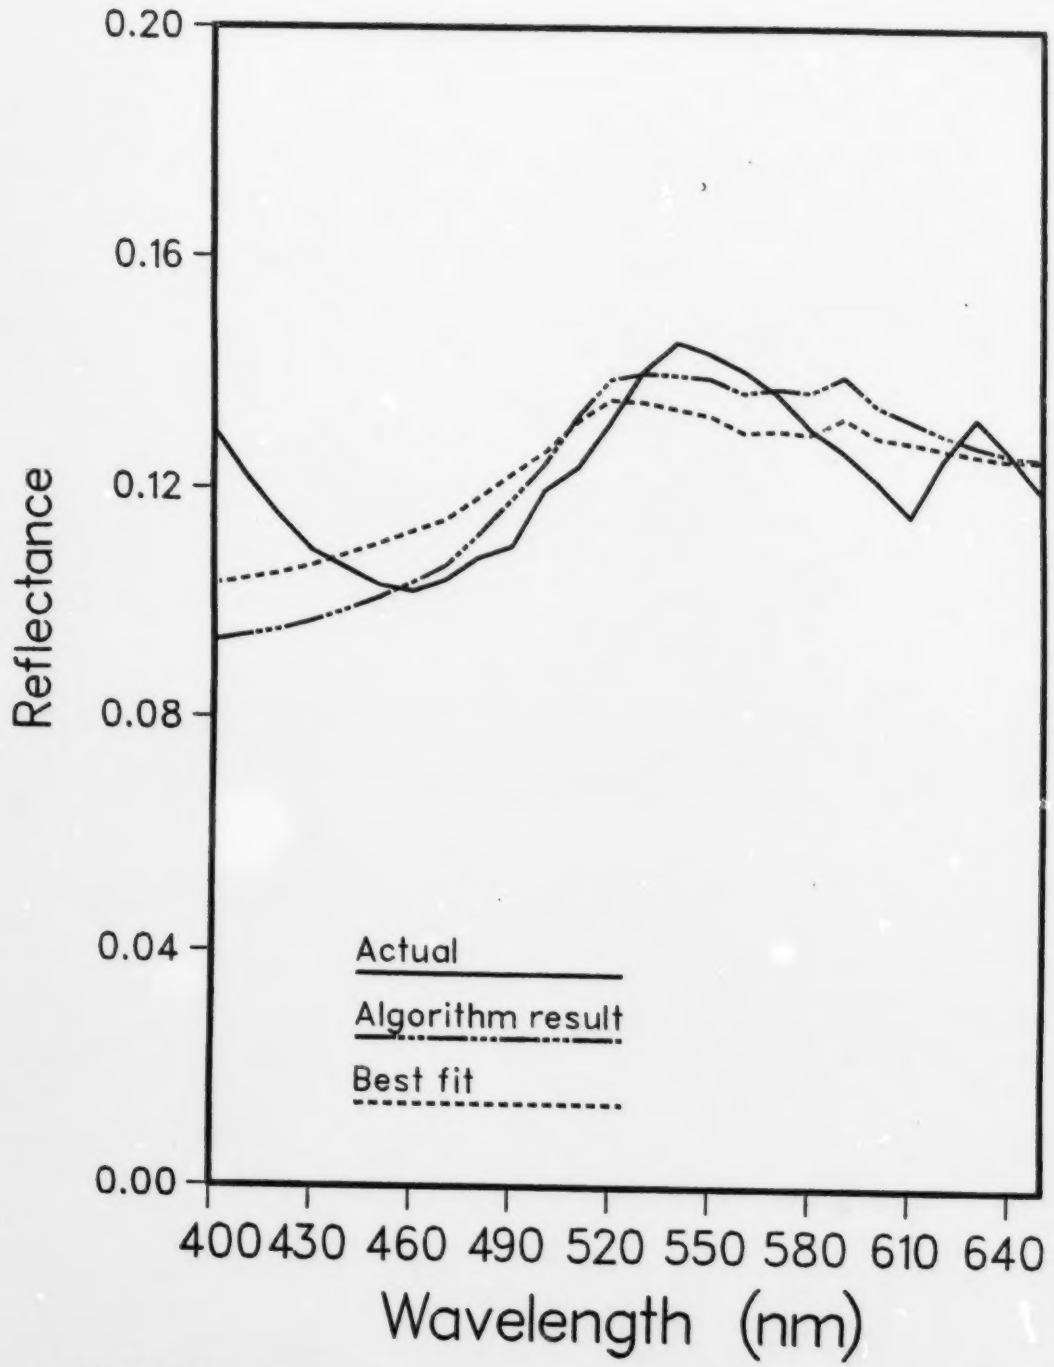

Fig.3. Comparison of the best fit given by dimension 3 finite dimensional model with surface spectral curve resulting from the algorithm. The best fit error (here, the reflectance curve is Krinov #54) has error=8.11%, whereas the algorithm result has error=9.75%, a factor of 1.202 worse than the best fit. When the geometrical mutual reflection factor α is 0.05, the algorithm gives the value 0.04983, so that the ratio of the estimate to the actual α is 0.9966

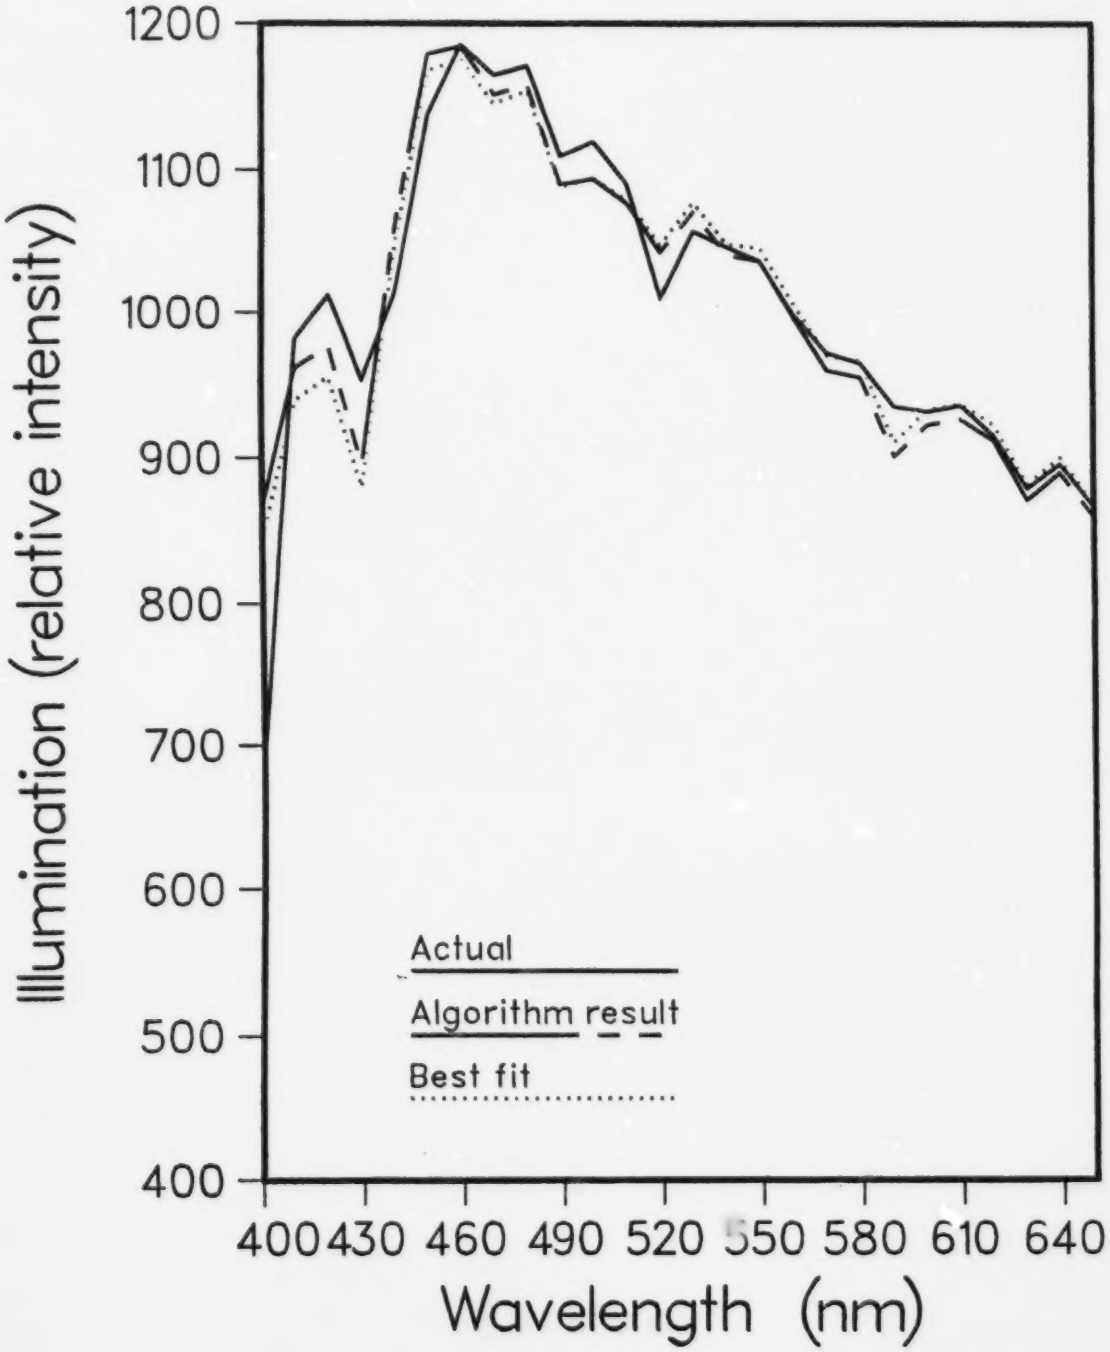

Fig.4. The original Dixon illumination spectrum for Australian daylight and results of the mutual reflection algorithm applied to the color signal of Fig.2 — error 4.13%.